

09/29/00

Jc922 U.S. PTO

10-02-00

A

IBM Docket No. FR9-1999-0080 US1

**In the United States Patent and Trademark Office
Patent Application Transmittal****Transmitted herewith for filing is the Patent Application of:**

Inventors(s): Ange Aznar, Daniel Orsatti, Jean-Claude Zunino, Bruno Rousseau

For: System and Method for Enabling Remote Surveillance of ATM Network Switching Node
Ports**Enclosed are**

31 pages of specification, including 11 claims, plus 7 sheets of drawings.

X An assignment of the invention to International Business Machines Corporation, Armonk, New York 10504.

X A certified copy of a/an *European* application.

X Declaration and Power of Attorney.

PTO-1449 & references

X A return post card

Other:

Filing Fee Calculation (For Other Than Small Entity)

Basic Fee:						\$690.00
Claims Fees:	Filed	Limit	Extra		Rate per Extra	
Total claims:	11	20	0		\$18.00	\$0.00
Independent claims:	2	3	0		\$78.00	\$0.00
	Multiple Dependent Claim Presented				\$260.00	\$0.00
	Total					\$690.00

Please charge Deposit Account **09-0464** for the **Total** set forth above. The Commissioner is authorized to charge payment of any additional filing fees required under 37 CFR §1.16 and any patent application processing fees under 37 CFR §1.17 or to credit any overpayment to the identified account. A duplicate copy of this sheet is enclosed.

Express Mail Certificate

Express Mail Label No: EI855835001US

Date: Sept. 29, 2000

I hereby certify that I am depositing the papers identified above with the U.S. Postal Service "Express Mail Post Office to Address" service on the above date, addressed to the Assistant Commissioner for Patents, Washington, DC 20231

Sandra S. Christopher

Sandra S. Christopher

BY: *Daniel E. McConnell*

Daniel E. McConnell

Attorney of Record Reg. No. 20,360

Date: Sept. 29, 2000

IBM Corporation 2Y7/B656
Intellectual Property Law
PO Box 12195
Res. Tri. Park, NC 27709

Telephone: 919-543- 1105 FAX 919-543-3634

SYSTEM AND METHOD FOR ENABLING REMOTE SURVEILLANCE OF ATM NETWORK SWITCHING NODE PORTS

BACKGROUND OF THE INVENTION

1. Technical Field:

The present invention relates generally to ATM (Asynchronous Transfer Mode) multi-node networks and more specifically to a system and method for enabling the surveillance, through the nodes and switching pieces of equipment forming the network, of any switching node port from any convenient remote access point.

2. Description of the Related Art:

Broadband ISDN (Integrated Services Digital Network) was an attempt to set up a single unified, worldwide high-speed network in place of the multiplicity of existing networks for different applications. On one hand, the new, universal network was intended to replace the functions of current speech, data and television networks and, on the other hand, to provide enough scope for the implementation of future communications technologies.

The first work on standards for the universal network of the future was begun by CCITT (International Telegraph and Telephone Consultative Committee) in the late 1980s, under the title "B-ISDN". B-ISDN is based on ATM which is a data transmission technique belonging to the family of cell switched systems (cell relay). Unlike

packet-switched systems, in which data packets of
variable length can be multiplexed over a line interface,
the length of a cell relay data packet is fixed and
simply referred to as a cell. ATM is a specific
5 implementation of cell relay and an integral element of
the CCITT specification for B-ISDN.

ATM, as the name suggests, utilizes an asynchronous
time division multiplexing scheme so that the data
10 streams to be transmitted are converted into fixed-size
cells and transferred asynchronously over a same physical
medium between network nodes. The allocation of the
units of information to the different transmission
channels is carried out using numerical channel
15 identifiers attached to each cell under the form of a VPi
(Virtual Path identifier) and a VCi (Virtual Channel
identifier).

ATM is a cell switching technology in which
20 cell-switch units within network nodes route cells
towards their final destination and which may also buffer
the cells before transmission to a next node. Cell
switching implies that a temporary end-to-end
transmission path must be provided before any data can be
25 transferred between sender and receiver. Thus,
cell-switched networks can take advantage of the greatly
varying bandwidth requirements in data communications by
allocating unused transmission capacity to other virtual
connections on an as needed basis.

30 Since cell-switched networks create only virtual
transmission links (over a common physical transmission

line), it is possible for cells to be lost if the memory capacity of the switching nodes is exceeded. This transfer procedure, which first requires a (virtual) link to be set up between the users, is said to be connection-oriented (as opposed to connectionless like e.g., IP the Internet Protocol). Links are categorized as either PVC (Permanent Virtual Connection), in which the link is set up once for all, as SVC (Switched Virtual Connection) in which the link is set up dynamically, on demand.

It is only after such a connection has been established, be it permanent or dynamic, that sending station can forward its data, under the form of cells, to the receiving station which receives them in the same order. However, if problems occur during the transfer of the data such as a buffer overload, resulting in discarded cells or the receipt of faulty or misdirected cells, the problem is immediately reported to the other end, which can react appropriately (generally by repeating the transmission of the cells forming a message so that the problem can be handled by the error recovery routines of the higher level protocols).

Thus, ATM switching units are key elements in any B-ISDN network. The fact that all ATM cells are the same size is exploited to implement multiport switching fabrics (typically 16x16 or 32x32) based on various efficient architectures which, when combined with the latest sub-micron fabrication processes, permit very high performances and aggregate throughputs which must be expressed in terabits per sec (10^{12} /sec). The remaining

task of the switching fabric and switches at each node of a network then becomes to provide transmission paths between the input ports and output ports requested at any given time in such a way that the fewest possible conflicts occur.

For example, an internal conflict to a switch fabric may occur if two or more cells are competing for the same output port at the same time. Despite every effort being made during the architecture and design phases of switches and cell-networks in general to provide the necessary mechanisms to handle the data flows without conflicts, traffic congestion is inevitable in a node when either overall traffic is increasing or when a traffic spike occurs at the node at a given instant. Also, flaws and shortcomings that may be present in some of the numerous hardware and software pieces that together implement a network, although not serious enough to prevent the network from operating reasonably well and, in any case, at the satisfaction of the end users, may also trigger occasionally some of the previously mentioned problems such as cell discarding.

Therefore, it is of utmost importance for those in charge of running and maintaining such an ATM network to be able to watch it and analyze it, if not constantly, at least any time it is necessary, from any convenient access point, without having to disturb user traffic whatsoever. Moreover, because ATM networks are typically used nowadays to implement high performance backbones of many medium or large networks (contradicting the initial expectation that ATM would become the universal means to

transport data, voice and all multimedia information, up to desktop) quality performance must be ensured with the proper tools, such that the high level of quality assurance expected of an ATM network is indeed delivered.

FR9-1999-0080

SUMMARY OF THE INVENTION

A method and a system enabling remote surveillance of any entry port to any switching node of an ATM cell-relay network are disclosed. Such an ATM cell-relay network usually comprises several switching nodes each having a plurality of ports. The method first assumes that a path is set up, through the ATM network, between an entry port and a remote observation point. Then, all cells of the incoming traffic, entering the entry port to be watched, are duplicated and marked. Subsequently, the cells are transported, following the path, up to the observation point. The preferred method of marking all duplicated cells comprises reserving one bit of the cells channel identifier to unambiguously distinguish them from the regular cells. This bit is asserted in every duplicated cell, such that the cells remains unaltered and are processed expeditiously in the intermediate switching nodes en route to the observation point. The method and system of the invention permit any entry port of an ATM network that may well span over large geographic areas, be conveniently observable and analyzed transparently i.e., without disturbing users traffic, from a remote location, such that the network can be maintained and operated trouble free.

All objects, features, and advantages of the present invention will become apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 depicts a switch within an ATM network;

Figure 2 illustrates an exemplary ATM network in which a preferred embodiment of the present invention may be practiced;

Figure 3 depicts a switch including a switch fabric with adaptive blades interfacing data communications lines in accordance with a preferred embodiment of the present invention;

Figure 4 is a block diagram illustrating cells replication within either the switch fabric or in blades in accordance with a preferred embodiment of the present invention;

Figure 5 depicts an ATM cell structure assembled in accordance with a preferred embodiment of the present invention;

Figure 6 is a logic flow diagram illustrating the process by which a cell enters a blade in accordance with

a preferred embodiment of the present invention; and

Figure 7 is a logic flow diagram depicting the process by which a cell exits the switch fabric to a blade in accordance with a preferred embodiment of the present invention.

5

FR9-1999-0080

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

With reference now to the figures and in particular with reference to **Figure 1**, there is illustrated an ATM switch **100**. In this figure, as an example, a 16x16 port switch **100** is shown. Each of the 16 ports, such as port **110**, is a bi-directional port capable of receiving and sending ATM cells **120** carrying a connection identifier under the form of a VPi (Virtual Patch identifier) **121** and a VCi (Virtual Channel identifier) **122** part of the header **125** of every cell.

Switch **100** performs two basic jobs. First, it identifies the connection identifier of every incoming cell from the VPi and VCi fields mentioned hereinabove. Switch **100** then transports cells from input ports to output ports (e.g. **130**) along a predetermined path towards their final destinations. The actual transport is carried out via a switching fabric **140** that provides dynamic transmission paths between the input ports and the output ports requested at any given time in such a way that the fewest possible conflicts occur like when two cells, from two different input ports, are competing for the same output port at the same time.

There are numerous architectures that have been proposed and implemented which all tend to reduce the transmission delay while increasing the aggregate throughput that can be handled through a switch. Nowadays, the latter parameter must be expressed in terabits/sec or 10^{12} bits/second. At the same time,

switch designs tend to drastically limit the number of cells that must be discarded due to conflicts of the kind mentioned above. Cell loss rates as low as 10^{-10} are reported, meaning that less than one cell out 10^{10} cells going through a switch need to be discarded under nominal conditions. Therefore, ATM switches indeed provide the high level of performance required to address the exponentially growing demand for bandwidth from transport demands by a variety of traffic such as pure data, voice, still images and video on a unified network.

Figure 2 shows a high performance ATM backbone network **220** comprising five switches including switches **201** and **202**. Switches are interconnected through high speed links such as link **210**, which in the depicted example, may be optical lines conforming to the well known SDH (Synchronous Digital Hierarchy) or SONET (Synchronous Optical Network) transmission standards respectively defined by the telecommunication standardization sector of the International Telecommunications Union or ITU-T in various G-series standards and by the American National Standards Institute or ANSI e.g., in T1.105. SDH and SONET have transmission rates that are compatible at 155.2 Mbps (referred to as STM-1 and STS-3 for SDH and SONET respectively), 622 Mbps (STM-4 and STS-12) and 2.48 Gbps (STM-16 and STS-48).

Very high performance backbone networks, such as backbone network **220**, are capable of moving huge amounts of data over large distances that may well span over states and countries. Backbone network **220** also permit

interconnection of thousands to millions of end-users including those on a Token Ring LAN **230**, which in the depicted example is located in northern Europe. Token ring LAN **230** may require connection with southern European users on an Ethernet LAN **240**, while both LANs may need to access resources from a third remote location **250**.

It is of the utmost importance that backbone network **220** be operational 24-hour a day, 7-day a week. Various redundancy-based techniques are known in the art for achieving near 100% availability. Hence, it is often the case that inter-switch transmission lines are actually duplicated as shown in a duplicated link **211**, such that a "hot" standby link is always ready to replace a failing link. This type of network must be closely managed and cannot be left unattended due to the enormity of traffic depending on its proper performance.

The present invention provides a method and system for monitoring the network from any access point to backbone network **220**. From any such access point, it should be possible to monitor a remote switch port **231** (or any other entry port to any of the switching nodes forming the network **220** irrespective of their location at the periphery of, or within backbone network **220**) from a distant location **260** where measuring equipment has been installed on a port **261** of switch **201** for the purpose of snooping all traffic entering port **231** or exiting port **231** from the group of users connected on Token Ring **230**. This monitoring must be achievable through a direct link

210 between switches 201 and 202, or through any alternate path 215 encompassing other switches of the backbone if more appropriate.

5 The present invention permits such monitoring without any additional hardware, aside from tooling equipment in distant location 260. In addition the present invention eliminates the need for installation of dedicated lines between switches. This latter
10 requirement is often automatically fulfilled in backbone network in which high speed links are duplicated for the above-mentioned reason, thus providing spare bandwidth that can be utilized occasionally for the surveillance of the network.

15 **Figure 3** illustrates a closer view of a few ports of a switch fabric 300. Each active port of a switch is equipped with a port adapter hereinafter referred to as an adaptive blade or simply a "blade". Three such blades
20 are shown as blades 310, 320 and 330, that connect to three bi-directional switch fabric ports 301, 302 and 303, respectively. **Figure 3** further depicts blades interface data communications lines, such as line 317 in a variety of combinations including the number of lines, their speeds, the standards they are complying to and the
25 types of physical interface e.g., optical or electrical.

30 A typical example of blade interface line 317 could be a blade interfacing sixteen optical STM-1/STS-3 communications lines 315 at 155.2 Mbps. Such a blade interface is compatible with the European SDH and US

SONET standards. Blades must be able to accommodate the incoming and outgoing traffic of the sixteen lines to a single switch-fabric port and thus must be capable of handling an average bi-directional traffic up to 2.4 Gbps.

The role of a blade such as blade **310** is thus to concentrate and dispatch traffic from/to several lines from/to a higher speed switch fabric port so as to expand, as much as possible, the connectivity of switch fabric **300** and take full advantage of its intrinsic performance capabilities. A switching node is therefore comprised of a switching fabric having ports to which blades are connected to allow adaptation to one or more transmission lines which become, in turn, ports of the switching node. The term "port", as employed in the following description, refers either to the switching fabric ports themselves or to the ports of the switching node (i.e., the end points of the transmission lines connected to the switch blades). The distinction in the intended meaning of "port" is not ambiguous from within the context. The connectivity of a switching node may be further expanded by utilizing 155 Mbps line **317** to link to or from more than one local area network through any convenient peripheral box as shown in **Figure 2**.

A blade includes a receive side **311** that handles all the incoming traffic received from the sixteen lines, and a transmit side **312** that handles all the outgoing traffic to be dispatched over the sixteen lines **315** from the switch-fabric port.

The chief function of receive side **311** is to inspect each received cell, thus enforcing policing and performing the verification of the connections according to the Quality of Service (QoS) defined for each connection and possibly tagging or discarding nonconforming cells. The main task of receive side **311** is to determine the destination of an incoming cell (i.e., the port and blade it should be sent to) such that the cell follows the path assigned by the network manager (ATM is a connection oriented protocol which assumes that a path exists before traffic can start).

The assigned path is found from within a connections table **316** that tracks all active virtual connections supported at any given time by the blade. Therefore, whenever a cell is received, table **316** is interrogated to determine through which outgoing port the received cell must exit. Table **316** appends to the incoming cell the necessary information such that the cell can find its way out through the switch fabric. Depending on the switch fabric design and implementation, there are many ways in which this objective may be achieved. Irrespective of the particular method chosen, **Figure 3** shows incoming cells that are input to switch fabric **300** through port **301** to which blade **310** is connected. Switch fabric **300** responds by forwarding the incoming cells to the transmit side of the appropriate target blade.

At this point, it is worth noting that more than one switch fabric output port may need to be specified in the information appended to the cell because some cells must be forwarded to several ports for multicasting purposes.

This function is generally supported within switch fabric 300 which is capable of replicating the incoming cells such that they are dispatched over several ports as necessary. Moreover, another level of multicasting may have to be performed in the transmit side of the blade. For example, in multi-line blade 320, similar to blade 310, a cell exiting switch fabric port 302 may have to be replicated several times so that the cell may be dispatched over the sixteen lines that blade 320 interfaces or a multiple thereof.

To illustrate such cell duplication, two cases of duplication are shown in **Figure 3**. A first duplication that occurs within switch fabric 300 shows that the incoming flow of cells 318 may have to be duplicated over ports 302 and 304. Flow 318 may need to be duplicated again, on the transmit side of blade 320 (i.e., over lines 321 and 324) when multicasting requires that destinations accessible through those lines need to receive the multicast flow of cells. This information is found through the interrogation of a local connection table 326 for each cell received through the switch fabric port 302 in a manner similar to what is done for all incoming cells through line 317 and table 316.

Figure 4 starts describing a preferred embodiment of the invention which is broadly referred to as port snooping in the remaining description. Snooping a node port (i.e., of all the incoming traffic arriving through a line such as 417), so that the traffic can be remotely observed, first consists in real-time marking of all the

cells entering the blade from the port. Then a copy of them is transported up to an observation point, possibly through many intermediate other nodes. Upon entering snooping mode, a snoop source point of interest is chosen from which a path must be setup up to the observation point. As with all regular ATM connections the path is set up in advance and must have a reserved bandwidth compatible with the actual throughput of the snooped port so that its traffic can be safely transported and observed from a remote location.

This path setup and bandwidth reservation was previously illustrated in **Figure 2** where path **215** traverses five switching nodes before reaching the, thus enabling observation point **260** to receive a copy of all incoming traffic passing through port **231** and that originating from Token Ring LAN **230**. Set up of this path is done under the control of the network manager through the control point of each ATM switch with standard methods and techniques known from the art.

What is different from the prior art is that the control point of the switching node is instructed, by the overall network manager or whatever combination of tools and people are responsible for keeping the network up and running, to mark all cells arriving to port that must be snooped. Techniques for marking the cells are further discussed in the following. Whatever method is used, all cells (irrespective of the channel identifier they are carrying) of an incoming flow **418** from line **417**, that is chosen to be snooped, are, upon receipt, marked **419** in blade **410**. Cells of flow **418** may subsequently have to be

5 duplicated in the switch fabric itself (using the same overall mechanism as used for multicasting) so that they are sent to the observation point through another blade 430. An extra flow of cells 438 is thus created which is the image of the snooped flow 418, and which follows the path reserved in the network for it up to the last blade where measuring and observation equipment is connected.

10 This duplication of the snooped flow may also have to be performed on the transmit side of a blade whenever the snooped flow and regular connections use the same blade for exiting a particular switching node. This case is illustrated in **Figure 4** with blade 420. Because of the multicast function that must be supported at
15 switching nodes (discussed hereinabove with reference to **Figure 3**), cells belonging to some connections listed in table 426 may also have to be replicated locally multiple times (e.g., on lines 421 and 424). Therefore, snooped flow 428 must be further replicated in addition to the
20 replication required by the multicasting of cells belonging to regular connections, if any is required. Consequently, it is possible to send replicated marked cells over the same line as the line borrowed by a regular connection. This is illustrated **Figure 4** where
25 snooped flow goes through line 421.

30 The above described generation and handling of the marked cells needs to be carried out only in the source node where a port is snooped. All the other nodes, on the path to the observation point, will automatically recognize and handle the marked cells to their final destination.

Figure 5 depicts the manner in which cells can be marked so that they are recognized as snooped traffic and are automatically routed, unaltered, to their final destination according to a preferred embodiment of the present invention. ATM cell header **500** is 5-bytes long including a one-byte CRC (called HEC) **510** for transmission protection. The remaining portion of header **500** carries four bytes of information including a 16-bit VCI (Virtual Channel identifier) **520** and a 8-bit or 12-bit VPI (Virtual Path identifier) **530**, depending upon the interface and is said to be UNI (User to Network Interface) **540** or NNI (Network to Network Interface) **550**. Therefore, at least 24 remaining bits are available within the header to identify a connection.

Because, unlike other protocols, this is not a universal address or identifier, it has only a local significance between two switches and can freely be assigned by the network manager. There is no danger of a shortage of identifiers to differentiate the virtual connections even on the highest speed lines. The only constraint is that identifiers must be unique on the same physical communications line. The simplest method for marking the cells is to borrow one bit of the VPI or VCI fields which, when turned on, signifies that the cell belongs to the snooped traffic and must be processed accordingly. Alternatively, on the UNI interface, which includes a 4-bit GFC (Generic Flow Control) field **560**, one of the bits could be devoted to marking the cells. Thus, in a preferred embodiment of the invention, the MSB (Most Significant Bit i.e., bit 15) of the VCI field **570** is chosen to mark the snooped cells. Hence, it becomes

network manager responsibility not to use this bit for channel identifiers for the regular connections.

Figure 6 illustrates a method for processing cells on the receive side of every blade in accordance with a preferred embodiment of the present invention. For each incoming cell **600**, the header is tested **605** to determine if cell has already been marked in a previous node at the origin of the snooped flow. If answer is positive then, next step **610** is to check if the current blade is indeed on the path of the snooped flow as defined by the network manager. If the answer is negative an error situation is encountered and cell must be discarded at step **611**.

Normally, if snoop path has been correctly defined the answer to question **610** is positive in which case the marked cell can be forwarded directly to the switch fabric as shown at step **630** after the necessary information has been added to it, so the cell can be steered through the switch fabric and reach the output port from which it will be transported to a next node along the snoop path. What is appended to the cell to cross the switch fabric is highly dependent on the switch fabric design. However, it is generally broadly referred to as a 'Switch Tag' **615**. This latter comprises the necessary information to direct the incoming cell to the right port i.e., the outgoing snoop port **620** in this case. Thus, the processing of marked cells, arriving at intermediate nodes, is expedited and those cells are moved along the snoop path unaltered until they reach the observation point.

However, if incoming cell 600 is not marked so that the answer to question 605 is negative, it is a regular cell (non-snoop cell) possibly belonging to one of the virtual connection recognized on the receive side of the current blade. Through the interrogation of a connection table 316 (shown in **Figure 3**) a determination is made of whether the cell is actually expected because it corresponds to a predefined connection. If the answer to question 625 is positive (the normal case) then, the cell follows the regular process i.e., a switch tag is added 615 which specifies the outgoing connection port 635 so that the cell eventually traverses the switch fabric and reaches the output port following the connection it belongs to.

However, if current blade has been set, by the network manager, as being the source of the snoop flow, snoop traffic must be originated from this blade. Therefore, if answer to the next question 640 is positive the switch tag must be altered so that the switch fabric is instructed to replicate the regular cell to the snoop port 645 also, eventually resulting in the forwarding of a marked cell. If current blade is not a snoop source, the answer to question 640 is negative and the above branch is just bypassed. At this stage it is important to remark that the snoop port and connection port may, for some of the incoming cells, match. In such a case the replication of the regular cell is not handled by the switch fabric itself but rather by the transmit side of the outgoing blade. This is further discussed with reference to **Figure 7**.

Finally, if answer to question **625** is negative, the cell is not recognized and should normally be discarded **650**. However, if blade has been declared as the source of the snoop flow, so that answer to question **655** is positive, all cells arriving at the snooped port entry must be forwarded to the observation point requiring that a switch tag be added **615** specifying the snoop port **620** like previously described for the incoming marked cells.

Figure 7 is a logic flow diagram depicting the process by which a cell exits the switch fabric to a blade in accordance with a preferred embodiment of the present invention. The method depicted in **Figure 7** focuses on the transmit side of a blade. Each regular cell received from a switch fabric port triggers the interrogation of local connection table **326** of **Figure 3**. The objective of this interrogation is many fold. First, a determination is made as to which output a virtual connection cell should be sent. As previously mentioned, this determination may result in several or possible all lines being selected if the cell must be multicast. Therefore, a multicast mechanism allowing replication of cells must be supported in the transmit side of every blade.

A VCi/VPi swap (a standard operation in ATM switches) is also performed as shown at step **750**. Since, in ATM networks, channel identifiers have only a local significance, VPi and VCi fields of cell headers are possibly exchanged at each node even though cells are following predetermined virtual paths. Numerous methods of achieving these functions are known the art, each

differing according to the actual implementation of switches and blades. Irrespective of such differences of any particular implementation, all methods include the step of adding a correlator to the switch tag wherein the switch tag is appended to the cells received from the lines in the receive side of each blade. This correlator functions as an index that is used to address the local connection table found in each blade transmit side. In conjunction with the local connection table, the correlator is utilized to obtain information for handling the current incoming cell.

Included within the information obtained by utilizing the correlator with the local connection table is the identity of output port(s) that the current cell should be sent and also what VPi/VCi is to be used for each of them. This determination is generally performed within a linked list of destinations which is traversed each time a new cell is received so that the cell may be replicated accordingly. If no multicast is required (i.e., cell is sent on a single outgoing port) then the list becomes a one-item list. Snooping in accordance with the principles set for herein, advantageously utilizes features of multicasting which is merged into the process whenever necessary as explained hereafter.

The algorithm for processing a cell that is received from a switch fabric port on the transmit side of every blade, first includes the step of testing a bit in the switch tag that was appended to the cell to determine whether the cell must be 'recirculated' (i.e., must be replicated), such that the cell participates in the

generation of the snoop flow.

It should be noted that not all switch fabric are, per se, "ATM switches" but rather are cell-switches. If the switched cells are truly ATM cells, sufficient information will be appended to guide the cell within the switch fabric itself and the blades. Such appended information is broadly referred to herein as a switch tag as previously discussed. A switch tag includes information imparted by the receive side for the purpose of multicasting of a cell on more than one switch fabric port. The switch tag also includes the above-mentioned correlator that is used as an entry point into the local connection table.

Thus, among the overhead bits added to the ATM cell which is to be switched, one bit is reserved for permitting a snoop recirculation of cells by the transmit side of the blades when necessary. The first step is then to test if this bit is set or not as shown at step 700. When a cell is received from the switch fabric this bit cannot be active since this bit is not set by either the receive side or the switch fabric. In response to a negative response at inquiry step 700, the method proceeds to the interrogation of the local connection table (step 710) to determine whether the cell is expected as being a part of a predetermined flow. If answer is negative, the next step 720 is to determine whether the cell is marked as belonging to the snoop flow. If answer is negative, an error case is detected and logged, and the cell is discarded as shown at step 730.

However, if answer is positive, a snoop cell (marked in current node or a previous node) is recognized and forwarded to the snoop port. This branch of the algorithm corresponds to what was depicted in **Figure 4** and flow **438**. If answer to question **710** was however, positive in which case a cell belonging to a registered connection is to be processed the next step **750** is the VPi/VCi swap, a standard operation in ATM network, previously discussed. As shown at next question **760**, a determination is made whether the current cell must be replicated. One way of achieving this, as already previously suggested, is to organize in the local connection table linked list of connections on which cells have to be replicated. Hence, unicast and multicast are handled alike i.e., when the end of the linked list is reached (which may comprise only one item) replication is stopped.

If one assumes first that answer to step **760** is negative, a multicast recirculation bit is set **770** in the overhead bits of the cell (i.e., the switch tag) after which cell is forwarded **780** through the selected output port. Obviously, when cell is eventually sent over a telecommunications line, to which output port is tied, so as it is transported to another node, all the overhead (the switch tag) that was appended to the ATM cell for allowing internal processing, is removed. At this stage algorithm resumes in **700**. When, at first loop or during a subsequent loop (if cell is multicast), the end of the linked list is eventually reached and the answer to question **760** is positive then, the next step **790** is to wonder if a snoop cell must be generated too. If this is

indeed the case the snoop recirculation bit is set 795.
Hence, when algorithm resume in 700 answer is finally
positive so as branch ending in step 740 is gone through
resulting in the forwarding of a snoop cell.

While the invention has been particularly shown and
described with reference to a preferred embodiment, it
will be understood by those skilled in the art that
various changes in form and detail may be made therein
without departing from the spirit and scope of the
invention.

CLAIMS

What is claimed is:

1. A method of enabling remote surveillance of any entry port to any switching node of an ATM cell-relay network, said ATM cell-relay network comprising one or more of said switching nodes that each include a plurality of ports, said method comprising the steps of:

setting up a path, through said ATM network, between said entry port and a remote observation point;

duplicating all cells of incoming traffic entering through said entry port;

marking all of said duplicated cells; and

transporting, along said path, said marked and duplicated cells up to said observation point.

2. The method of claim 1, wherein said step of marking all duplicated cells further includes the steps of:

reserving one bit within a cells channel identifier field to unambiguously distinguish said duplicated cells; and

asserting said reserved bit of every of said duplicated cells.

1 3. The method of claim 2, wherein said step of
2 reserving one bit further comprises the step of
3 alternatively utilizing one bit of a cells Generic Flow
4 Control field.

1 4. The method of claim 2, wherein said step of
2 reserving one bit further comprises the step of reserving
3 a most significant bit of a Virtual Channel identifier
4 field part of said cells channel identifier field.

1 5. The method of claim 1, wherein said marking step
2 further comprises the step of marking said duplicated
3 cells with any unique combination of bits within a header
4 of said duplicated cells.

1 6. The method of claim 1, wherein said step of
2 duplicating cells is performed within said switching node
3 of said entry port.

1 7. The method of claim 1, wherein said step of
2 duplicating cells further comprises the step of
3 duplicating unexpected cells, errored cells, and
4 nonconforming cells.

1 8. The method of claim 1, wherein within all
2 intermediate said switching nodes along said path up to
3 said observation point, said step of transporting said
4 marked and duplicated cells further comprises the steps
5 of:

6 recognizing said marked and duplicated cells upon
7 entry of each of said marked and duplicated cells into a

8 port of each intermediate switching node;

9 testing to determine whether said port of each
10 intermediate switching node is included within said path
11 to said remote observation point;

12 in response to determining that said port is
13 included within said path to said remote observation
14 point:

15 maintaining said marked and duplicated cells
16 unaltered; and

17 moving said marked and duplicated cell forward;

18 in response to determining that said port is not
19 included within said path to said remote observation
20 point:

21 discarding said marked and duplicated cells; and

22 recording an error.

1 9. A system for enabling remote surveillance of any
2 entry port to any switching node of an ATM cell-relay
3 network, said ATM cell-relay network comprising one or
4 more of said switching nodes that each include a
5 plurality of ports, said system comprising:

6
7 means for setting up a path, through said ATM
8 network, between said entry port and a remote observation
9 point;

10 means for duplicating all cells of incoming traffic
11 entering through said entry port;

12 means for marking all of said duplicated cells; and

13 means for transporting, along said path, said marked
14 and duplicated cells up to said observation point.

1 10. The system of claim 9, further comprising:

2 in each of said switching node:

3 a cell switch fabric; and

4 adaptive blades between said ports and said cell
5 switch fabric.
6

1 11. The system of claim 10, further comprising means
2 within said cell switch fabric or within said adaptive
3 blades for replicating cells.

ABSTRACT OF THE DISCLOSURE
SYSTEM AND METHOD FOR ENABLING REMOTE SURVEILLANCE OF ATM
NETWORK SWITCHING NODE PORTS

5 In an ATM cell-relay network usually comprising
several switching nodes a method enabling remote
surveillance of any entry port to any switching node of
the network. The method first assumes that a path is set
up from the entry port to a remote observation point.
10 Then, all cells of the incoming traffic, entering entry
port to be watched, are duplicated and marked. After
which they are transported, unaltered, following the
path, up to the observation point. The invention permits
that any entry port of an ATM network, that may well span
15 over large geographic areas, be conveniently observable
and analyzed transparently i.e., without disturbing users
traffic, from a remote location so as network can be
maintained and run trouble free.




Figure 1

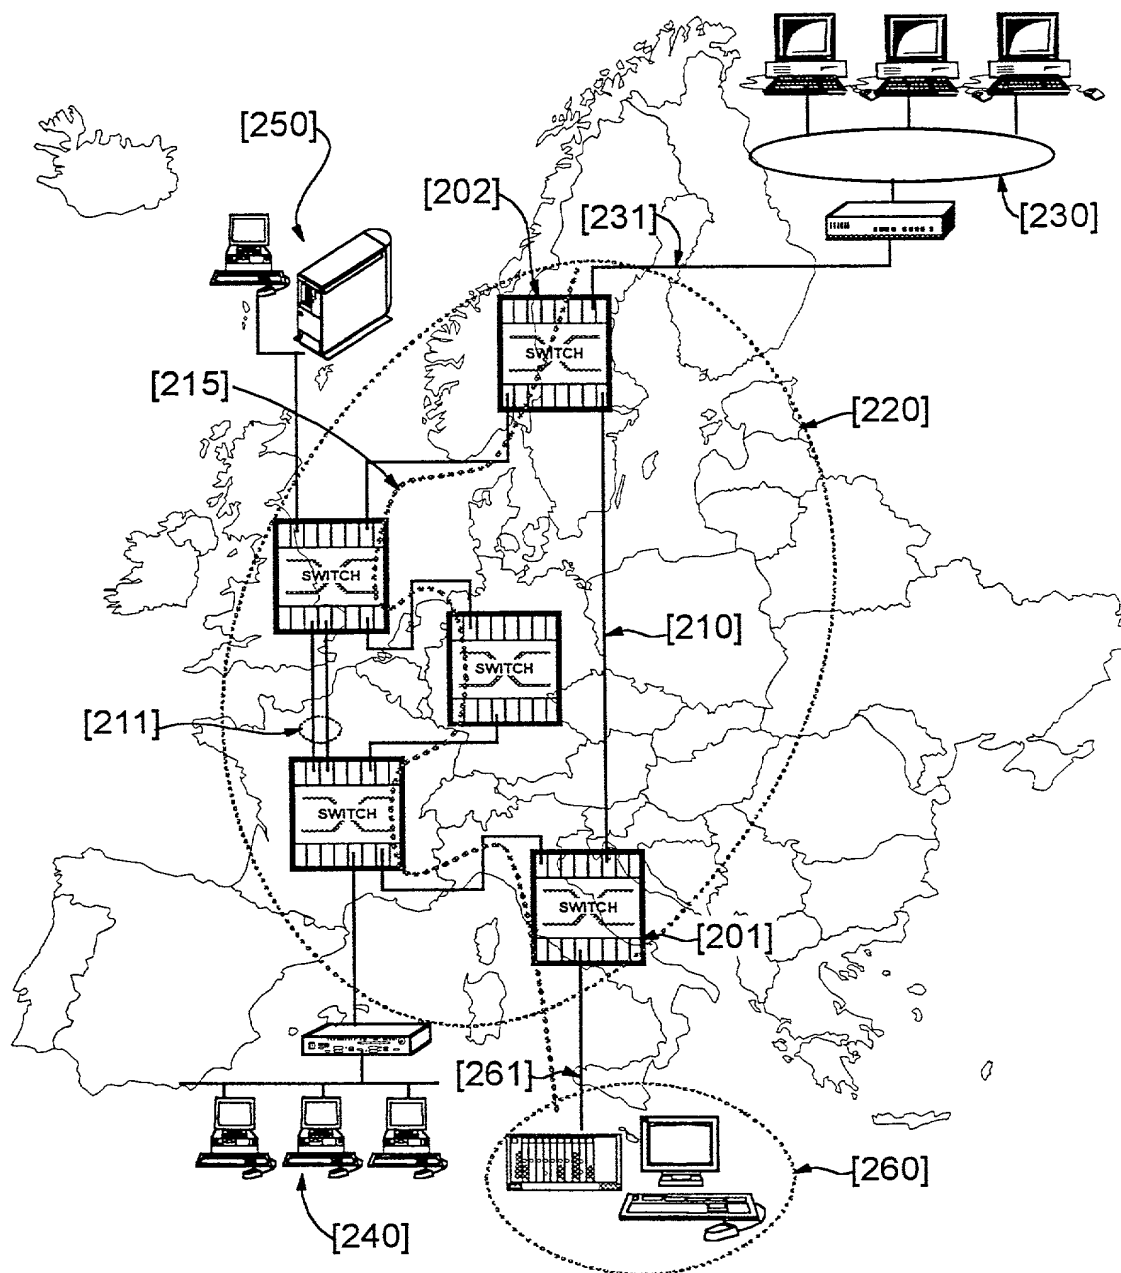



Figure 2



[311]

[312]

[326]

[318]

[316]

[315]

x16

[317]

[321]

[324]

Figure 3

Figure 4




Figure 5

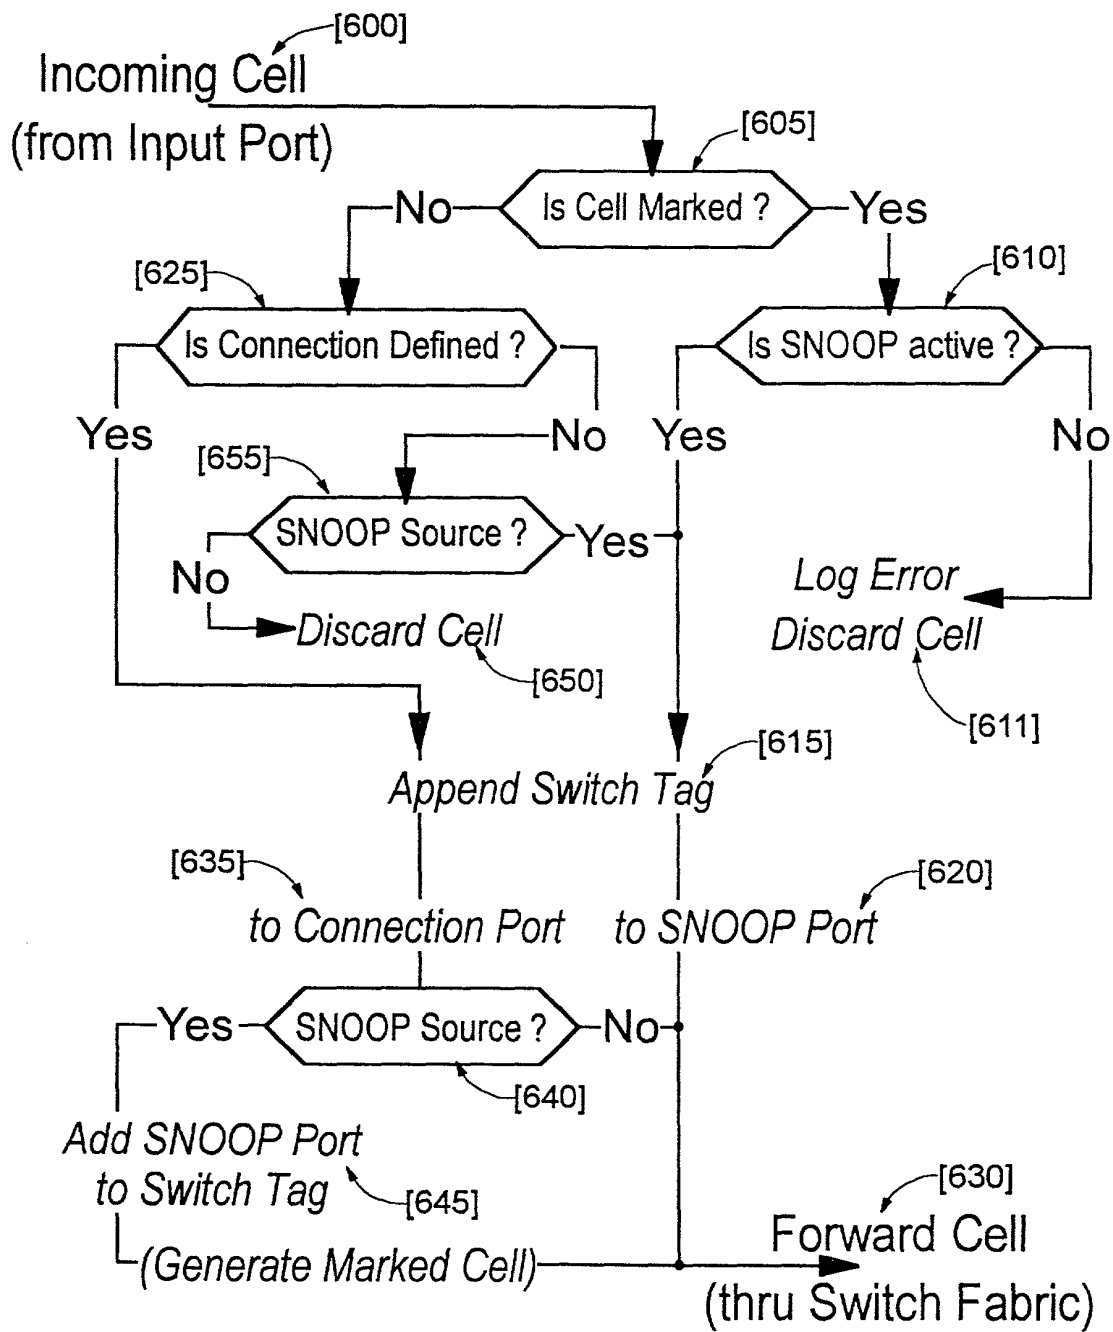


Figure 6

Incoming Cell
(from Switch Fabric Port)

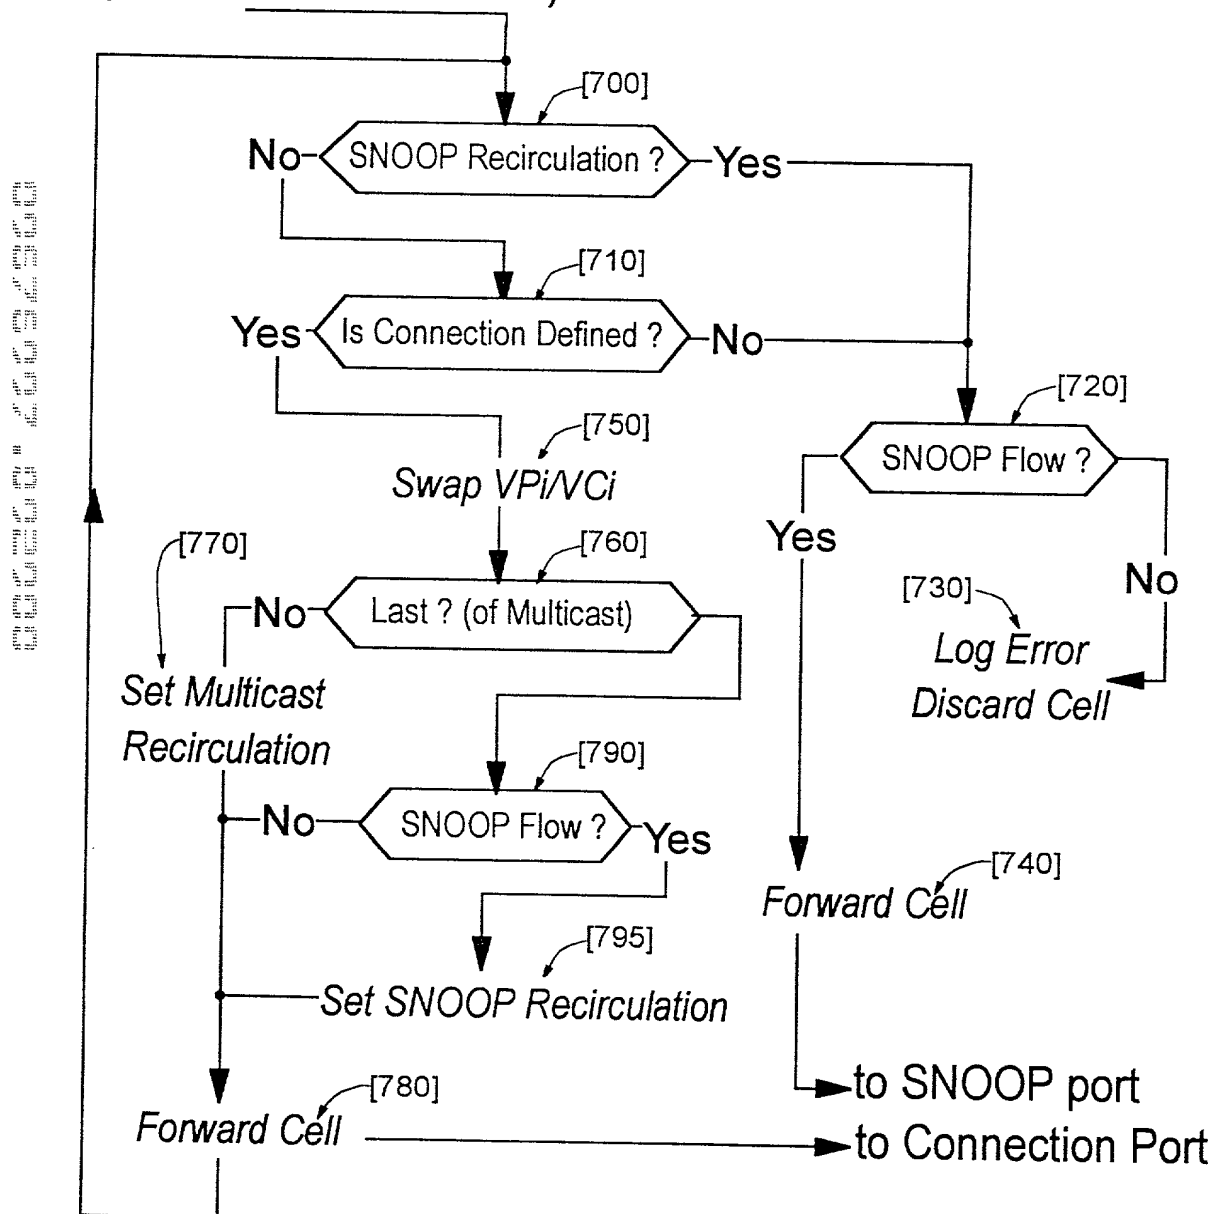


Figure 7

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name; I believe I am an original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled:

**System and Method for Enabling Remote Surveillance of ATM
Network Switching Node Ports**

the specification of which is identified by the attorney (IBM) Docket Number appearing above.

I hereby state that I have reviewed and understand the contents of the above- identified specification, including the claims.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

<u>Number</u>	<u>Country</u>	<u>Day/Month/Year</u>	<u>Priority Claimed</u>
99480110.8	Europe	21/10/99	Yes

I hereby claim the benefit (a) under Title 35, United States Code, §119(e) of any U.S. application listed below and identified as a provisional application or (b) under Title 35, United States Code, §120 of any U.S. application listed below and not identified as a provisional application, and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior U.S. application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application

Prior U.S. Applications

<u>Serial No.</u>	<u>Filing Date</u>	<u>Status</u>
-------------------	--------------------	---------------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Page 1 of &

As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith: Daniel E. McConnell, Reg. No.

IBM Docket No. FR9-1999-0080 US1

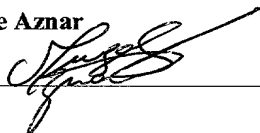
20,360; Joscelyn G. Cockburn, Reg. No. 27,069; Horace St. Julian, Reg. No. 30,329; and Christopher A. Hughes, Reg. No. 26,914.

Send all correspondence to: Daniel E. McConnell, IBM Corporation 2Y7/B656; PO Box 12195; Research Triangle Park, NC 27709.

First Inventor:

Ange Aznar

Signature:



09/13/2000

Date

Residence: 1612 Chemin de l'Escours
06480 La Colle s/Loup - France

Citizenship: French

Post Office Address: Same as above

Second Inventor:

Daniel Orsatti

Signature:



31/8/2000

Date

Residence: 900 Chemin des Colles et Regagnades
06610 La Gaude - France

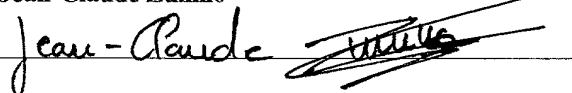
Citizenship: French

Post Office Address: Same as above

Third Inventor:

Jean-Claude Zunino

Signature:



09/13/2000

Date

Residence: Jardins Corniche Fleuri 141
06200 Nice - France

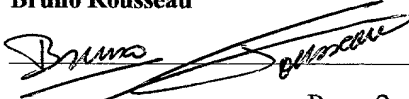
Citizenship: French

Post Office Address: Same as above

Fourth Inventor:

Bruno Rousseau

Signature:



20/8/2000

Date

Residence: Residence Les Chenes Verts

IBM Docket No. FR9-1999-0080 US1

327 av. General Leclerc
06140 Vence - France

Citizenship: French

Post Office Address: Same as above